

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 833 479 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:
16.03.2005 Bulletin 2005/11

(51) Int Cl.7: **H04L 12/407**, H04L 12/28

(21) Application number: **96307200.4**

(22) Date of filing: **27.09.1996**

(54) **Contention resolution process for data networks**

Konkurrenzbetriebsauflösungsverfahren für Datennetzwerke

Procédé de résolution de contention pour des réseaux de données

(84) Designated Contracting States:
DE FI FR GB SE

(43) Date of publication of application:
01.04.1998 Bulletin 1998/14

(73) Proprietor: **Hewlett-Packard Company,
A Delaware Corporation
Palo Alto, CA 94304 (US)**

(72) Inventors:
• **Spratt, Michael Peter
Bath BA1 2XN (GB)**
• **Sharon, Oran
Tivon 36056 (IL)**

(74) Representative: **Coker, David Graeme et al
Hewlett-Packard Limited
Intellectual Property Section
Building 2
Filton Road
Stoke Gifford Bristol BS12 6QZ (GB)**

(56) References cited:
EP-A- 0 140 077 US-A- 4 860 006
US-A- 5 392 033

- **PHOENIX CONFERENCE ON COMPUTERS AND COMMUNICATIONS, 20 - 22 March 1985, SCOTTSDALE, ARIZONA, pages 404-412, XP002025435 M.E.ULUG: "Analysis of the Preemptive Contention Law"**
- **ELEKTRONIK, vol. 31, no. 8, April 1982, MÜNCHEN, DEUTSCHLAND, pages 67-70, XP002025436 HUBERT EING, HANS J. WILHELMY: "Das SP-Netz - ein neues, vielversprechendes lokales Netzwerk (LAN)"**

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 0 833 479 B1

Description

[0001] This invention relates to data networks.

[0002] In the last 20 years, Local Area Networks (LANs) have revolutionised the way in which computers have been used in the work place, and have allowed personal computers ("PC"s) and the client/server paradigm to replace to a large extent the mainframe/terminal technology prior to LANs. In a similar way, in the last 10 years, the MIDI network has revolutionised the production and performance of music, and created a huge market in MIDI-compatible equipment.

[0003] In a similar way, a home LAN could help to revolutionise home equipment, by bringing together computing and electronic entertainment. Such a network might link around the home digital TVs, a set-top box, digital VCR, CD player producing digital audio, PC, printer and ISDN router. Currently, there are multiple standards bodies examining home LANs, but no really strong contender for home LAN technology has emerged, and most proposals seem to have some drawbacks. The proposals tend to concentrate on carrying just video traffic and are not compatible with the Ethernet interfaces which are already installed or being installed in large numbers in the home, to connect PCs to printers and ISDN routers. These proposals ignore the fact that Ethernet is very widespread.

[0004] The present invention, or at least certain embodiments of it, is concerned with enabling Ethernet compatible video transmission, allowing the same network to be used for piping video around the home as the transport of computer data, so that existing Ethernet cards of the home network can be used without modification. The present invention, or at least certain embodiments of it, achieves this by enabling video and the like to be transmitted on the Ethernet with higher priority. The present invention, or at least certain embodiments of it, is also concerned with allowing video to be transmitted with very low jitter. This is particularly important in interfaces to home appliances, where the extremely cost sensitive market means only minimal buffering at the receiver can be afforded to overcome network jitter. The present invention, or at least certain embodiments of it, is also concerned with providing a protocol which is simple, and/or implementable purely in hardware, and/or robust.

[0005] US-A-5392033 (Oman et al.) describes a networking system comprising a central arbiter which uses priority generators to determine which resource (station) on a shared bus has priority in accessing the bus. A number of protocols for assigning priorities are described on the basis of first-come first-served, round robin or a combination of these ranking methods.

[0006] US-A-4860006 (Barall) relates to a collision avoidance method for use in networks with a plurality of stations. The method allocates unique transmission slot times to each station in reference to a heartbeat signal which is generated independently of a data transmis-

sion.

[0007] EP-A-0140077 also describes a collision avoidance protocol for use in a network with a plurality of stations. This method provides for a network arbitration period following the termination of each transmitted frame which is divided into a plurality of bus access windows which are then assigned to each respective station in the network.

[0008] In accordance with a first aspect of the present invention, there is provided a method of operating a data network having a bus and a plurality of stations adapted to transmit data on the bus, wherein:

the stations have respective rankings and perform a contention resolution process whereby they transmit in cycles and within each cycle they transmit in an order determined from their rankings;

wherein, the contention resolution process performed by each station prior to transmission of data on the bus comprises the steps of:

transmitting a jam on the bus for up to a period of time dependent on the current ranking of that station;
detecting whether there is a collision on the bus while transmitting the jam; and
if not, enabling transmission of data; and

each station is adapted to perform a fairness process, which comprises the steps of:

detecting whether one of the stations has transmitted data succeeding a jam without a collision;

if not, disabling its contention resolution process until such a detection has been made; and
if a station has further data to transmit immediately after it has already transmitted data in a current cycle said station refrains from the contention resolution process until it has determined that the current cycle has been completed.

[0009] Thus, this aspect of the invention allows the stations to transmit in turn. It will be appreciated from the following description that not all of the stations on the network are necessarily arranged to operate in accordance with this aspect of the invention, but only what will hereinafter be referred to as "high-priority" stations, such as those transmitting digital video and audio.

[0010] It should be noted that in this specification, unless the context otherwise requires, "bus" is not limited to a bus employing a cable connection, and also includes any kind of shared communications channel, such as wireless buses employing, for example, radio or infra-red as a means of communication.

[0011] Preferably, the contention resolution process performed by each station prior to transmission of data on the bus comprises the steps of: transmitting a jam on

the bus for up to a period of time dependent on the current ranking of that station; detecting whether there is a collision on the bus while transmitting the jam; and, if not, enabling transmission of data; and each station is operable to perform a fairness process, which comprises the steps of: detecting whether one of the stations has transmitted data succeeding a jam without a collision; and, if not, disabling its contention resolution process until such a detection has been made. This therefore provides an easily implementable protocol for contention resolution and fairness.

[0012] In order to deal with the possibility that a station may have taken part unsuccessfully in contention resolution, but then some change occurs which means that it no longer requires to transmit data, for example if that station is switched off, the fairness process of each station preferably comprises the further steps of: detecting whether its contention resolution process has been disabled for a predetermined period of time; and if so, re-enabling its contention resolution process.

[0013] The ranking of each station may be fixed, and this may provide reasonably efficient contention resolution if all of the stations are active all or most of the time. However, in order to allow the efficiency of the contention resolution process to adapt as stations become active and inactive, each station preferably performs a ranking allocation process prior to an initial transmission on the bus, which comprises the steps of: determining, from the bus, rankings currently in use by the other station(s); and allocating to itself a current ranking to one extreme of the determined rankings. In this case, the ranking allocation process performed by each station preferably comprises the further steps of: determining whether its current ranking is the same as that of the other station, or another of the stations; and, if so, recommencing its ranking allocation process. In order to deal with the problem that, if the error comes to light at the time of contention resolution, the rankings reallocated by the stations may also be the same, preferably the method further comprises the step of waiting for a pseudo-random period of time before the recommencement of the ranking allocation process. In order to reduce the delays which take place during contention resolution for stations with particular (e.g. "high") rankings, the ranking allocation process performed by each station preferably comprises the further steps of: detecting whether the ranking next to that station's current ranking in the direction towards the other extreme of the rankings (e.g. the lower direction) continues to be in use; and, if not, re-allocating to itself that other ranking as its current ranking.

[0014] A second aspect of the present invention provides a data network comprising a bus and a plurality of stations operable to transmit data on the bus, the stations being arranged to perform a method according to the first aspect of the invention.

[0015] A third aspect of the present invention provides a station adapted to form part of a network according to

the second aspect of the invention.

[0016] As mentioned above, the stations may be thought of as high-priority stations, and the method of the first aspect of the invention may also involve at least one low-priority station which is operable to transmit data on the bus, the low-priority station(s) performing a different contention resolution process (for example in accordance with IEEE standard 802.3 CSMA/CD) to the high-priority stations.

[0017] In the case where the or each low-priority station is operable to perform a contention resolution process prior to transmission of data on the bus, the contention resolution process commencing when that station does not sense any activity on the bus, or a predetermined period of time after that station senses a transition from carrier to silence on the bus, or a predetermined period of time after that station has finished transmitting data, each high-priority station is preferably operable, immediately after transmission of data, to transmit a filler on the bus for a period of time exceeding said predetermined period(s) of time for the low-priority station(s).

[0018] This latter feature may be employed in methods other than in accordance with the first aspect of the invention. Therefore, according to a fourth aspect of the present invention, there is provided a method of operating a computer network having a bus, and a plurality of high-priority stations and at least one low-priority station each operable to transmit data on the bus, wherein: the or each low-priority station is operable to perform a contention resolution process prior to transmission of data on the bus (for example in accordance with IEEE standard 802.3 CSMA/CD), the contention resolution process commencing when that station does not sense any activity on the bus, or a predetermined period of time after that station senses a transition from carrier to silence on the bus, or a predetermined period of time after that station has transmitted data; and each high-priority station is operable, immediately after transmission of data, to transmit a filler on the bus for a period of time exceeding said predetermined period(s) of time for the low-priority station(s).

[0019] Preferably, the or each low-priority station ceases transmission, at least temporarily, after transmission of data.

[0020] A specific embodiment of the present invention will now be described by way of example with reference to the accompanying drawings, in which:

Figure 1 illustrates a network having high-priority and low-priority stations;

Figure 2 is a flow diagram showing a tag allocation and re-allocation process performed by high-priority stations on the network;

Figure 3 is a timing diagram illustrating contention resolution between high-priority stations on the network;

Figure 4 comprising Figures 4A and 4B is a flow diagram showing a contention resolution process performed by each of the high-priority stations on the network; and

Figure 5 is a time-space diagram illustrating contention resolution on the network after the transmission of a frame by a low-priority station, with time advancing in the x direction and position along the network bus shown in the y direction.

[0021] Referring to Figure 1, a network comprises a plurality of high-priority stations 10a, 10b, 10c, 10d and a plurality of low-priority stations 12a, 12b, 12c, connected by Ethernet cable 14. The low-priority stations 12a etc. might include an IBM-compatible PC, a printer and an ISDN router having conventional network hardware and software and operating in a conventional manner in accordance with the IEEE 802.3 CSMA/CD standard. The high-priority stations 10a etc. might include a digital television set, a digital TV set-top box, a digital video recorder/player and a CD player producing digital audio.

[0022] The stations 10a, 10b, 10c, 10d are termed "high-priority" stations because they employ a contention resolution protocol which gives them priority on the network over the low-priority stations 12a etc. The contention resolution protocol of the high-priority stations also ensures fairness amongst themselves.

[0023] The high-priority contention resolution protocol employs "tags", which are numbers 1, 2, 3 and which appear as a field in a high-priority frame. The allocation of tags by each high-priority station will now be described with reference to Figure 2. When a high-priority station wishes to initiate a session/connection to the network, in steps 20 and 22 it observes the network for a predetermined observation period such as 100ms and observes the tags, if any, which are currently in use by the other high-priority stations, and records the highest tag ("RHT") which is in use. If no other high-priority stations are active, then an RHT of zero is recorded. In step 24, the high-priority station allocates to itself a tag T which is one higher than the recorded highest tag RHT. The observation period of 100ms prevents a high-priority station possibly losing its tag place in the order to another high-priority station, provided that it transmits every 100ms or less. This would typically be the case if video or audio were transmitted at high priority. If necessary the figure of 100ms could be adjusted upwards, for example to accommodate stream-type applications which could have inter-transmission gaps of more than 100ms.

[0024] Once allocated, a high-priority station's tag T may change. As shown by steps 26 and 28, if the station is transmitting and it detects a tag collision, the process returns to step 20, after a pseudo-random delay provided in step 30, and the station's tag is re-allocated from scratch. A tag collision might be detected, in this case,

if the station reaches the end of a long jam (described below) while still detecting a collision, and then stops transmitting but observes that no other high-priority station continues by sending a frame. The pseudo-random delay is introduced by step 30 in order to reduce the chances that the stations whose tags have collided will both allocate themselves the same tag in the subsequent steps 20 to 24.

[0025] As shown by step 26, the station's tag is also re-allocated from scratch if it is not transmitting and it observes that another high-priority station is using its tag T. In this case, there is no need for a pseudo-random delay, because the other station can continue without re-allocating its tag.

[0026] As shown by steps 26, 32, if a station is not transmitting, and it observes that no other high-priority station has transmitted within a predetermined preceding period, such as 100ms, with a tag T-1 one less than its own tag T, then it decrements its own tag by one, ie T becomes T-1, but not if its previous tag was one. Thus, the tags tend to float down to fill the lowest available numbers from 1 upwards.

TABLE 1

Station	10a	10b	10c	10d
Initially	1	2	3	4
10b stops transmitting	1		3	4
100ms later	1		2	4
A further 100ms later	1		2	3

[0027] For example, as shown in Table 1, the high-priority stations 10a, 10b, 10c, 10d at one stage have tags 1, 2, 3, 4, respectively. Station 10b then stops transmitting frames, so that the active tags are 1, 3, 4 for stations 10a, 10c, 10d, respectively. About 100ms later, station 10c has observed that tag 2 is spare and re-allocates its tag from 3 to 2. Then, about 100ms after that, station 10d has observed that tag 3 is spare and re-allocates its tag from 4 to 3.

[0028] The tags which are allocated and reallocated in the manner described above are used in contention resolution between the high-priority stations. Figure 3 shows a situation where a high-priority station is completing transmission of a frame 70, and four high-priority stations with tags 1, 2, 3, 4 are contending for transmission.

[0029] At the end of its frame 70, the station which is transmitting transmits an end-of-frame delimiter field EFD 72. As shown in step 40 of Figure 4A, each of the high-priority stations wishing to transmit detects the EFD and in step 42 initiates a respective long jam timer with a value K.T, where T is the respective tag, and K is a system constant which is chosen in dependence upon the propagation delay from one end of the network to the other and the time to detect a collision. In particular, K is preferably equal to or slightly greater than $2.(\tau + \delta)$,

where τ is the one-way end-to-end propagation delay of the network and δ is the maximum time for a high-priority station to detect an important transition on the network such as between carrier and silence, or between collision and no collision. (A jam which is transmitted on the network for a period of $K \geq 2 \cdot (\tau + \delta)$ is referred to in this specification as a "long jam.")

[0030] At step 44, each contending station starts to transmit a jam. As will be seen from steps 46, 48, each of the contending stations will transmit a jam for no longer than its tag T multiplied by the basic long jam period K . If, at step 46, one of the stations which is transmitting a jam detects that there has not been a collision on the network bus, then in step 50 it sets its high-priority collision flag HPCF to false (for a reason to be explained below), and then in steps 52, 54 ceases to transmit the jam and proceeds with transmitting its preamble and frame. If, at step 46, the high-priority station detects that there has been a collision on the bus which has ceased, then in step 56 it sets its HPCF to true, and then proceeds with the steps 52, 54 of ceasing the jam and transmitting its preamble and frame.

[0031] From the above and referring in particular to Figure 3, it can be seen that the stations with tags 1, 2, 3, 4 start transmitting jams substantially simultaneously at step 44 when they each detect the end-of-frame delimiter EFD 72 for the previously transmitted frame 70, and there is therefore a collision 74 on the network. The station with tag 1 ceases transmitting its jam 76 a period K later at time 78, but the stations with tags 2, 3, 4 continue to cause the collision 74. The station with tag 2 ceases transmitting its jam 80 a further period K later at time 82, but the stations with tags 3, 4 continue to cause the collision 74. The station with tag 3 ceases transmitting its jam 84 a further period K later at time 86. The consequent cease of a collision on the network at time 90 is detected by the station with tag 4, and therefore as a result of steps 46, 56, 52, 54 (Figure 4), it ceases to transmit its jam 92 at time 94 and proceeds with transmitting its preamble and frame 96. After the frame has been transmitted, the high-priority station transmits the end-of-frame delimiter EFD, as shown in step 58 in Figure 4B and initiates its long jam timer in step 60 in preparation for the possibility of transmitting another high-priority frame.

[0032] Having described how the high-priority stations 10a-d contend for access to the network, there now follows a description of how the low-priority stations 12a-c are prevented from transmitting while the high-priority stations are persisting with transmissions. As mentioned above, the low-priority stations 10a-d operate in a conventional manner in accordance with the IEEE 802.3 CSMA/CD standard. That standard prescribes that a station can only start to transmit on the network in one of the following circumstances:

1. In a period when the station does not sense any activity on the network.

2. When an inter frame gap (IFG) expires after the station senses a transition from carrier to silence on the network. (The IFG is defined in the standard as being the period of 96 bits, i.e. 9.6 μ s.)

3. When an IFG expires after the station has transmitted a frame.

[0033] In order to ensure that a low-priority station does not start to contend for the network when a high-priority frame has been transmitted and other high-priority frames are waiting to be transmitted, the high-priority stations start contending for the network at time 44 in response to the EFD 72, but the station which has transmitted a frame continues after the EFD 72 to transmit a filler 98 for a predetermined period, which is less than a basic long jam period K but at least 2τ , as shown in steps 62, 64 in Figure 4B. Thus, the filler period is sufficiently long that, despite propagation delays on the bus, none of the stations will detect a period of silence on the bus. Therefore the low-priority stations do not have an opportunity to contend for access to the network.

[0034] Having described contention resolution between high and low-priority stations after transmission of a high-priority frame, there now follows a description, with reference to Figure 5 and the remainder of Figure 4A, of contention resolution between stations after transmission of a frame by a low-priority station. Unlike a high-priority frame, after transmission of a low-priority frame there silence and no EFD and no filler.

[0035] A high-priority station wishing to transmit, upon detecting the silence in step 40, initiates a respective low-priority collision detect timer in step 100 with a value of $IFG + P + J + 2 \cdot \tau + 3 \cdot \delta$, where IFG, τ and δ are as described above, and where P is the preamble period for a low-priority frame and J is the jam period for a low-priority frame. In step 102, the high-priority station then starts to transmit its preamble and frame. In step 104, the high-priority station observes the bus while it is transmitting. If the frame end is reached without a collision, then the high priority collision flag is set to false in step 106, and the process proceeds to step 58 in Figure 4B. If, however, a collision is detected in step 104 in Figure 4A while the preamble or frame is being transmitted, in step 108 the high-priority station starts to transmit a jam, and then in step 110 waits for the collision to end or for the low-priority collision detect timer to expire. If the collision ceases before the timer expires, this indicates that the collision was with one or more low-priority stations, and the high-priority station proceed to step 50 where its high-priority collision flag is set to false, and then to steps 52, 54 where the jam ceases and the high-priority preamble and frame are re-transmitted. On the other hand, if in step 110 the timer expires before the collision ceases, this indicates that the collision was caused at least in part by another high-priority station. The process therefore proceeds to step 42 and a con-

tention resolution process takes place between the high-priority stations, as described above.

[0036] Figure 5 shows the case where two low-priority stations 1 and 2 are situated at opposite ends of the network bus, and a high-priority station is situated near the low-priority station 1. The low-priority station 1 is finishing transmitting a frame at point B. The end of the frame reaches the other end of the bus at point D, a period τ later. The low-priority station 2 senses the transition from carrier to silence and starts its inter frame gap timer at point E, a period δ later. At point F, the inter frame gap timer of the low-priority station 2 expires, a period IFG later, and at point G, a period δ later the low-priority station 2 starts to transmit its preamble for a period P.

[0037] Meanwhile, the end of the low-priority frame passed the high-priority station at point H and silence began, and the high-priority station therefore has started to transmit a preamble and then a jam (not shown in Figure 5) as a result of steps 40, 102, 104, 108 in Figure 4A. Accordingly, there is a collision for the low-priority station 2 at point K at the end of its preamble, and therefore in accordance with the IEEE 802.3 CSMA/CD standard it may start to transmit a jam at point L, a period δ later, if there is a frame waiting. The jam of the low-priority station continues for a period J, whereupon the low-priority station ceases contending for the bus at point M, and the end of the jam reaches the other end of the bus at point N. Since the high-priority station is situated near the other end of the bus, it can therefore be seen from Figure 5 that in order for the high-priority station to succeed over the low-priority stations, the value set in its low-priority collision detect timer in step 100 of Figure 4A needs to be at least as great as $\tau + \delta + \tau + \text{IFG} + \delta + P + \delta + J$, or as mentioned above $\text{IFG} + P + J + 2\tau + 3\delta$.

[0038] In order to ensure fairness between the high-priority stations during contention resolution, there now follows a description of the purpose of the high-priority contention flags HPCF which were set to true or false in steps 50, 56, 106 in Figure 4A.

[0039] Referring to the remainder of Figure 4B, when a high-priority station, which has just transmitted a frame, has finished transmitting its filler in step 64, it determines in step 120 whether it has another frame waiting to be transmitted, and if so it determines in step 122 whether its HPCF is set to true or false. As mentioned above, if false, this indicates that the high-priority station did not contend with another high-priority station for access to the bus prior to transmitting its last frame, and in this case the process proceeds to step 44 in Figure 4A, where the high-priority station contends with any other high-priority stations which may, by now, require access to the bus. However, if the HPCF is true, then in step 124 the high-priority station initiates an error timer with a value greater than the time taken for a high-priority contention resolution process and transmission of a high-priority preamble and frame.

[0040] Subsequently, in step 126, the high-priority

station observes the bus and waits for an end of frame delimiter EFD or for the error timer to expire. If the timer expires, this indicates that the other high-priority station (s) which contended for access to the bus before the previous frame was transmitted no longer requires access (for example if it has been switched off). In this case, the process proceeds to step 40 in Figure 4A, where the high-priority station can start from scratch in contending for access to the bus. If, however, an EFD is detected in step 126, then in step 128 the high-priority station observes the bus and detects whether the length of any collision following the frame is at least as long as a basic long jam period K. If so, as will be appreciated from the description with reference to Figure 3, this indicates that two or more other high-priority stations are contending for access to the bus, and the process proceeds back to step 124 so that the station defers contending for access to the bus. However, if in step 128, the detected collision length is less than the basic long jam period K, this indicates that there is only one (or possibly no) other high-priority station contending for access to the bus, and the process proceeds to step 40 in Figure 4A so that the high-priority station can contend for access to the bus after that other high-priority station has transmitted its frame.

[0041] Accordingly, round robin fairness is provided. As an example, suppose that high-priority stations A to C have tags 1 to 3, respectively. When they first contend, A to C collide and C wins because it has the highest tag. C transmits its frame and sets its HPCF to true. Therefore C does not take part in the next contention resolution process, in which A and B collide, and B wins because it has the highest tag. B transmits its frame and sets its HPCF to true. Therefore both B and C do not take part in the next contention resolution process, in which A's jam is less than a basic long jam period K. A transmits its frame, and B and C (assuming they wish to transmit another frame) wait for the EFD of A's frame, so that at the end of that frame all of the stations A, B, C contend for access to the bus.

[0042] As explained previously, the primary purpose of the tag field is to enable resolution of contention between the high-priority stations. However, the tag field may also be used for other purposes. In the case where, the high-priority stations transmit with an Ethernet / IEEE 802.3 frame format, a sending station in general identifies itself with a unique source address. However, in the low-cost environment of consumer appliances, it may not be thought appropriate to assign individual unique source addresses to every station. Thus, for example, every Sony model III digital VCR might share the same source address A, and every Philips Mark 2 set top box might share the same address B. This is not a problem if the Sony Model III and the Philips Mark 2 are on the same network, but if there are two Philips Mark 2 models connected to the same network, there is an issue as to how a receiver would distinguish the frames from the two identical models with the two identical

source addresses. The tags which would, as described above, be allocated as being different on the two identical machines could therefore be used as a simple way of distinguishing between the machines.

[0043] Under the IEEE 802 protocol, each station has a 48-bit address. Preferably, the high-priority stations are allocated addresses in a particular range of the 48-bit addressing space, so that the low-priority stations are outside that range and never attempt to receive the high-priority packets on the network. Also the first few bits of each high-priority station's address preferably have a particular bit pattern so that the high-priority stations can also very quickly work out which are high-priority frames.

Claims

1. A method of operating a data network having a bus (14) and a plurality of stations (10a-d) adapted to transmit data on the bus, wherein:

the stations have respective rankings (T) and perform a contention resolution process whereby they transmit in cycles and within each cycle they transmit in an order determined from their rankings;

wherein, the contention resolution process performed by each station prior to transmission of data on the bus comprises the steps of:

transmitting (42, 44) a jam on the bus for up to a period of time dependent on the current ranking of that station;
detecting (46) whether there is a collision on the bus while transmitting the jam; and
if not, enabling transmission (54) of data; and

each station is adapted to perform a fairness process, which comprises the steps of:

detecting (46, 104) whether one of the stations has transmitted data succeeding a jam without a collision;

if not, disabling (126, 128) its contention resolution process until such a detection has been made; and
if a station has further data to transmit immediately after it has already transmitted data in a current cycle said station refrains from the contention resolution process until it has determined that the current cycle has been completed.

2. A method as claimed in claim 1, wherein the fairness process of each station comprises the further steps of:

detecting (126) whether its contention resolution process has been disabled for a predetermined period of time; and
if so, re-enabling its contention resolution process.

3. A method as claimed in any preceding claim, wherein the ranking of each station is fixed.

4. A method as claimed in either of claims 1 or 2, wherein each station performs a ranking allocation process prior to an initial transmission on the bus, which comprises the steps of:

determining (22), from the bus, rankings currently in use by the other station(s); and
allocating (24) to itself a current ranking to one extreme of the determined rankings.

5. A method as claimed in claim 4, wherein the ranking allocation process performed by each station comprises the further steps of:

determining (26, 28) whether its current ranking is the same as that of the other station, or another of the stations; and
if so, recommencing its ranking allocation process.

6. A method as claimed in claim 5, further comprising the step of waiting (30) for a pseudo-random period of time before the recommencement of the ranking allocation process.

7. A method as claimed in any of claims 4 to 6, wherein the ranking allocation process performed by each station comprises the further steps of:

detecting (26) whether the ranking next to that station's current ranking in the direction towards the other extreme of the rankings continues to be in use; and
if not, re-allocating to itself that other ranking as its current ranking.

8. A data network comprising a bus (14) and a plurality of stations (10a-d) adapted to transmit data on the bus, wherein said stations (10a-d) are allocated respective rankings and each station is further adapted to perform a contention resolution process according to any of claims 1-7.

Patentansprüche

1. Ein Verfahren zum Betreiben eines Datennetzwerks, das einen Bus (14) und eine Mehrzahl von Stationen (10a - d) aufweist, die angepasst sind, um

die Daten auf dem Bus zu senden, bei dem:

die Stationen jeweilige Rangordnungen (T) aufweisen und einen Konkurrenzauflösungsprozess durchführen, wodurch dieselben in Zyklen senden und dieselben innerhalb jedes Zyklus in einer Ordnung senden, die aus den Rangordnungen derselben bestimmt ist;

wobei der Konkurrenzauflösungsprozess, der durch jede Station vor einer Sendung von Daten auf dem Bus durchgeführt wird, folgende Schritte aufweist:

Senden (42, 44) einer Blockierung auf dem Bus für bis zu eine Zeitperiode, die von der aktuellen Rangordnung dieser Station abhängig ist;

Erfassen (46), ob es eine Kollision auf dem Bus gibt, während die Blockierung gesendet wird; und

falls nicht, Freigeben einer Sendung (54) von Daten; und

jede Station angepasst ist, um einen Gerechtigkeitsprozess durchzuführen, der folgende Schritte aufweist:

Erfassen (46, 104), ob eine der Stationen nach einer Blockierung ohne eine Kollision gesendet hat;

falls nicht, Sperren (126, 128) des Konkurrenzauflösungsprozesses derselben, bis eine derartige Erfassung vorgenommen wurde; und

falls eine Station weitere Daten aufweist, die gesendet werden sollen, unmittelbar nachdem dieselbe bereits Daten in einem aktuellen Zyklus gesendet hat, unterlässt die Station den Konkurrenzauflösungsprozess, bis dieselbe bestimmt hat, dass der aktuelle Zyklus abgeschlossen wurde.

2. Ein Verfahren gemäß Anspruch 1, bei dem der Gerechtigkeitsprozess jeder Station die folgenden weiteren Schritte aufweist:

Erfassen (126), ob der Konkurrenzauflösungsprozess derselben für eine vorbestimmte Zeitperiode gesperrt wurde; und

falls dem so ist, wieder Freigeben des Konkurrenzauflösungsprozesses derselben.

3. Ein Verfahren gemäß einem der vorhergehenden Ansprüche, bei dem die Rangordnung jeder Station

fest ist.

4. Ein Verfahren gemäß einem der Ansprüche 1 oder 2, bei dem jede Station vor einer anfänglichen Sendung auf dem Bus einen Rangordnungszuweisungsprozess durchführt, der folgende Schritte aufweist:

Bestimmen (22) von Rangordnungen von dem Bus, die durch die andere Station (die anderen Stationen) gegenwärtig in Verwendung sind; und

Zuweisen (24) einer aktuellen Rangordnung zu einem Extrem der bestimmten Rangordnungen für sich selbst.

5. Ein Verfahren gemäß Anspruch 4, bei dem der Rangordnungszuweisungsprozess, der durch jede Station durchgeführt wird, ferner folgende Schritte aufweist:

Bestimmen (26, 28), ob die aktuelle Rangordnung derselben die gleiche wie die der anderen Stationen oder einer anderen der Stationen ist; und

falls dem so ist, Empfehlen des Rangordnungszuweisungsprozesses derselben.

6. Ein Verfahren gemäß Anspruch 5, das ferner den Schritt eines Wartens (30) auf eine pseudozufällige Zeitperiode vor dem Wiederbeginn des Rangordnungszuweisungsprozesses aufweist.

7. Ein Verfahren gemäß einem der Ansprüche 4 bis 6, bei dem der Rangordnungszuweisungsprozess, der durch jede Station durchgeführt wird, ferner folgende Schritte aufweist:

Erfassen (26), ob die Rangordnung am nächsten zu der aktuellen Rangordnung dieser Station in die Richtung zu dem anderen Extrem der Rangordnungen hin in Verwendung bleibt; und

falls nicht, Neuzuteilen dieser anderen Rangordnung für sich selbst als die aktuelle Rangordnung derselben.

8. Ein Datennetzwerk, das einen Bus (14) und eine Mehrzahl von Stationen (10a - d) aufweist, die angepasst sind, um Daten auf dem Bus zu senden, wobei den Stationen (10a - d) jeweilige Rangordnungen zugewiesen sind und jede Station ferner angepasst ist, um einen Konkurrenzauflösungsprozess gemäß einem der Ansprüche 1 - 7 durchzuführen.

Revendications

1. Procédé de fonctionnement d'un réseau de données comportant un bus (14) et une pluralité de stations (10a-d) adaptées à transmettre des données sur le bus, dans lequel :

les stations possèdent des rangs respectifs (T) et exécutent un processus de résolution de conflit tel qu'elles transmettent dans des cycles et à l'intérieur de chaque cycle, elles transmettent dans un ordre déterminé par leurs rangs ;

dans lequel, le processus de résolution de conflit exécuté par chaque station avant transmission des données sur le bus comprend les étapes consistant à :

transmettre (42, 44) un blocage sur le bus pendant une période de temps allant jusqu'à une période dépendant du rang courant de cette station ;
détecter (46) s'il y a une collision sur le bus pendant la transmission du blocage ; et
si tel n'est pas le cas, autoriser la transmission (54) des données ; et
chaque station est adaptée à exécuter un processus d'équité, qui comprend les étapes consistant à :

détecter (46, 104) si l'une des stations a transmis des données à la suite d'un blocage sans collision ;
si tel n'est pas le cas, désactiver (126, 128) son processus de résolution de conflit jusqu'à ce qu'une telle détection ait été effectuée ; et
si une station a d'autres données à transmettre juste après avoir déjà transmis des données dans un cycle courant, ladite station n'exécute pas le processus de résolution de conflit jusqu'à ce qu'elle ait déterminé que le cycle courant est achevé.

2. Procédé selon la revendication 1, dans lequel le processus d'équité de chaque station comprend les autres étapes consistant à :

détecter (126) si son processus de résolution de conflit a été désactivé pendant une période de temps prédéterminée ; et
si tel est le cas, réactiver son processus de résolution de conflit.

3. Procédé selon l'une quelconque des revendications précédentes, dans lequel le rang de chaque station est fixe.

4. Procédé selon l'une ou l'autre des revendications 1 ou 2, dans lequel chaque station exécute un processus d'allocation de rang avant transmission initiale sur le bus, comprenant les étapes consistant à :

déterminer (22), à partir du bus, les rangs actuellement utilisés par la ou les autres stations ;
et
allouer (24) à elle-même un rang courant à une extrémité des rangs déterminés.

5. Procédé selon la revendication 4, dans lequel le processus d'allocation de rang exécuté par chaque station comprend les autres étapes consistant à :



déterminer (26, 28) si son rang courant est le même que celui de l'autre station, ou d'une autre des stations ; et
si tel est le cas, recommencer son processus d'allocation de rang.

6. Procédé selon la revendication 5, comprenant en outre l'étape consistant à attendre (30) pendant une période de temps pseudo-aléatoire avant le recommencement du processus d'allocation de rang.

7. Procédé selon l'une quelconque des revendications 4 à 6, dans lequel le processus d'allocation de rang exécuté par chaque station comprend les autres étapes consistant à :

détecter (26) si le rang qui suit le rang courant de cette station en se dirigeant vers l'autre extrémité des rangs continue à être utilisé ; et
si tel n'est pas le cas, réallouer à elle-même cet autre rang comme étant son rang courant.

8. Réseau de données comprenant un bus (14) et une pluralité de stations (10a-d) adaptées à transmettre des données sur le bus, dans lequel on alloue auxdites stations (10a-d) des rangs respectifs et chaque station est adaptée, en outre, à exécuter un processus de résolution de conflit selon l'une quelconque des revendications 1 à 7.



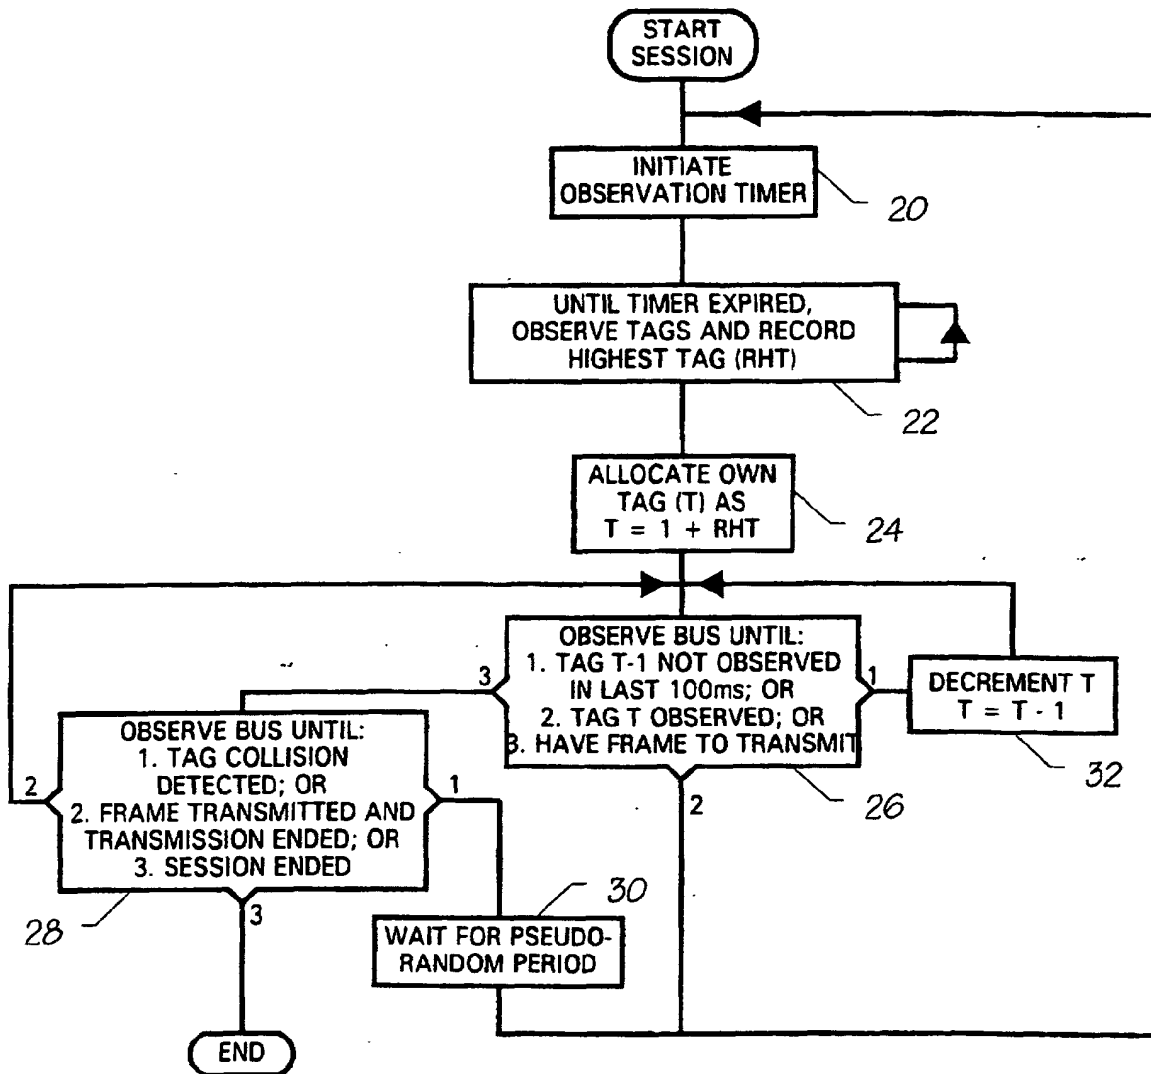


FIG. 2

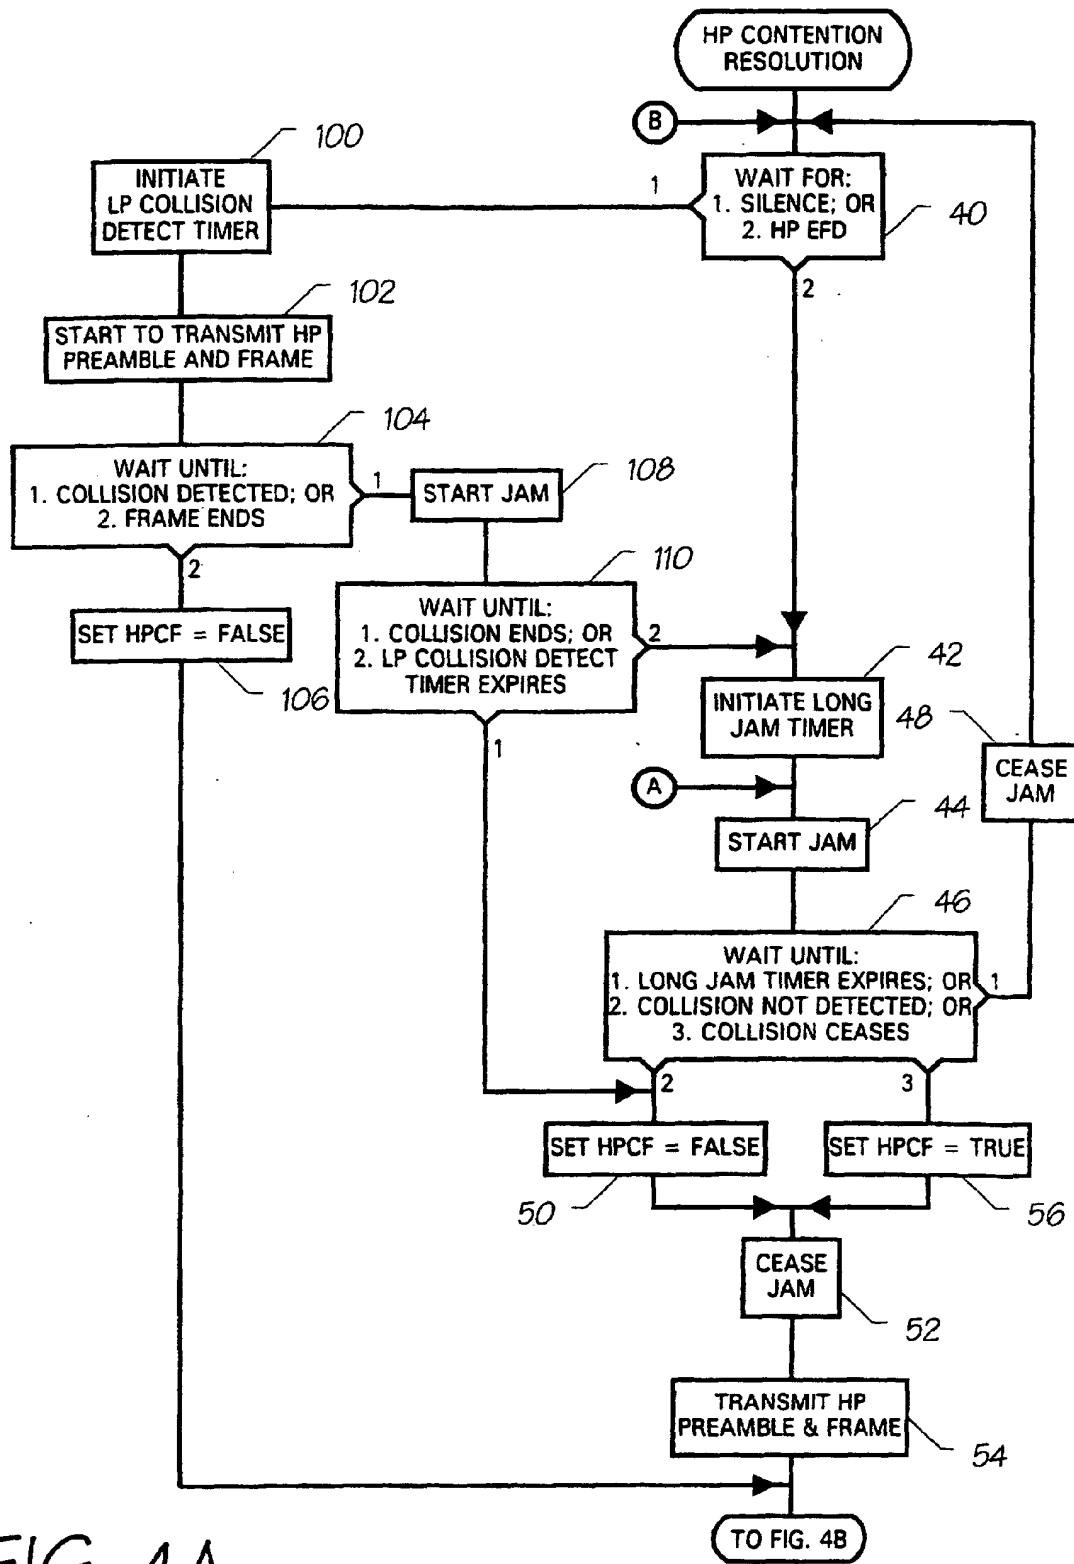


FIG. 4A




FIG. 4B




FIG. 5